

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: REVERSE DIVISION PROCESS

APPLICANT: PIERRE A. LAURENT

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EV 382041444 US

March 26, 2004
Date of Deposit

REVERSE DIVISION PROCESS

TECHNICAL FIELD

This patent application relates generally to a reverse
5 division process that is based on continued fraction
representations of quotients.

BACKGROUND

Any quotient of two numbers P and Q can be represented as
10 a continued fraction. A continued fraction representation of
the quotient of P/Q has the following form:

$$\begin{array}{rcl} P/Q = a_0 + & & 1 \\ & & \hline & & 1 \\ 15 & & a_1 + \hline & & 1 \\ & & a_2 + \hline & & 1 \\ & & a_3 + \hline 20 & & a_4 + \dots \end{array}$$

The coefficients a_1, a_2, a_3 , etc., are known as partial
25 quotients. The partial quotients of a continued fraction may
be represented in list notation, as follows:

$$P/Q = [a_0, a_1, a_2, a_3, a_4 \dots a(k)] \text{ for } k \geq 4.$$

Mathematics for generating a continued fraction is presented below, with reference to an example provided in prior art literature. In the following example, P has a value of 3796 and Q has a value of 1387. In this example, when 3796 is divided by 1387, the result is

$$3796 = 1387*2 + 1022,$$

where 1022 is the remainder. This equation can be rewritten as a continued fraction expansion, as follows:

$$3796/1387 = 2 + 1022/1387 = 2 + 1/(1387/1022).$$

Proceeding in the same manner as above,

$$1387/1022 = 1 + 365/1022 = 1 + 1/(1022/365).$$

Thus,

$$3796/1387 = 2 + 1022/1387 = 2 + 1/(1 + 1/(1022/365)).$$

Proceeding further in the same manner as above yields the following

$$1022 = 365*2 + 292$$

$$365 = 292*1 + 73$$

$$292 = 73*4.$$

Making the appropriate substitutions in the above continued fraction representation and rewriting the equations without the parentheses yields

$$\begin{aligned} 3796/1387 &= 2 + 1022/1387 = 2 + 1/1 + 1/(1022/365) = \\ &= 2 + 1/1 + 1/2 + 1/(365/292) \\ &= 2 + \cfrac{1}{1 + \cfrac{1}{2 + \cfrac{1}{1 + \cfrac{1}{4}}}} \end{aligned}$$

In list notation, the quotient of 3796/1387 is thus represented as [2, 1, 2, 1, 4].

DESCRIPTION OF THE DRAWINGS

Fig. 1 is a flowchart of a process for performing a reverse division process.

Fig. 2 is a block diagram of hardware that may be used to implement the reverse division process.

Fig. 3 is a graph showing signals generated via the reverse division process.

Fig. 4 is a graph showing a line between two points.

Fig. 5 shows the graph of Fig. 4 in which pixels selected via the reverse division process have been illuminated to generate the line.

5

DESCRIPTION

A reverse division process is described herein that uses continued fractions to obtain data. The reverse division process has broad applicability, but is particularly useful for processes that obtain and/or use periodic outputs, as described below. The mathematical basis for the reverse division process is illustrated as follows, using two integers P and Q.

15

The inverse of a continued fraction P/Q is defined by a sum of integer values $g(i)$ ($i \geq 1$) which fulfill the condition

$$g_1 + \dots + g(Q) = P.$$

For values of i greater than Q , $g(i)$ can be defined recursively to be

20

$$g(i) = g(i - Q).$$

25

The $g(i)$ values comprise elements of a periodic suite, $S(n)$, of numbers having a period of Q , the sum of which is defined as follows:

$$S(n) = g_1 + g_2 + \dots + g(n-1) + g(n), n \geq 2$$

$$S(0) = 0$$

5 The choice of $g(i)$ values is not unique, but there is one set of values $g(i)$ which provides a lowest division error in the continued fraction (caused, e.g., by remainders in the division process of rational numbers). This set of values satisfies, for any value n and m where $n \geq m + Q$,

10

$$| P - Q * ((S(n) - S(m)) / (n - m)) | < 1.$$

By way of example, for a quotient of $34/6$, a continued fraction representation is as follows:

15

$$P/Q = 34/6 = 5 + 1/(1 + 1/2), \text{ or}$$

$[5,1,2]$ in list notation. In this example, there are several sets of possible $g(i)$ values. For example:

20

$$6 + 6 + 6 + 6 + 5 + 5 = 34$$

may be used as $g(i)$ values. Other sets of $g(i)$ values include, but are not limited to

25

$$6 + 5 + 6 + 6 + 5 + 6 = 34$$

$$6 + 6 + 6 + 5 + 6 + 5 = 34$$

$$5 + 6 + 5 + 5 + 5 + 5 = 34$$

One way to reduce division error is to choose $g(i)$ values from the set $[a_0, a_0+1]$, and to interleave the $g(i)$ values

5 following the values of convergents $p(k)/q(k)$ (defined below) of the continued fraction P/Q .

More specifically, in the continued fraction

$$[a_1, a_2, a_3, \dots, a_k, \dots, a_n],$$

10 where $k \leq n$, the k th convergent ($C_k = p(k)/q(k)$) is the fraction that results from expanding the continued fraction to the partial quotient a_k . By way of example, for the continued fraction

15

$$P/Q = 1 + \frac{1}{2 + \frac{1}{3}}$$

20

three convergents, for $k=1, 2$ and 3 are

$$C_1 = 1$$

25

$$C_2 = 3/2$$

$$C_3 = 10/7$$

To reduce division remainder errors, $g(i)$ values, $g(n)$,
30 can be selected according to the following constraints:

$g(n) = a_0 + 1$ when $(S(n) - S(m)) / (n - m)$ is equal to an
even convergent $p(2k)/q(2k)$ and $g(m) = a_0 + 1$
 $g(n) = a_0$ otherwise

5

In the example give above, the following $g(i)$ values

$6 + 5 + 6 + 6 + 5 + 6$

satisfies the lowest error constraint.

10 The foregoing $g(i)$ selection process is approximated by
process 10 shown in Fig. 1. Process 10 may be implemented in
software or hardware. Fig. 2 shows hardware 12 that may be
used to implement process 10. Hardware 12 includes first
counters 14, second counters 16, and circuitry 18, such as a
15 controller, that controls the counters. In the case of
hardware, the counters comprise memory locations. Circuitry
18 controls inputs to, and outputs from, the counters.

 Prior to executing process 10, first counters 14 are
preloaded with partial quotients values $a_0 \dots a_n$ of a continued
20 fraction P/Q . The partial quotient values may be determined
manually or may be calculated using either hardware, software,
or a combination thereof. Each of first counters 14 has a
corresponding counter in second counters 16. Second counters
16 may be loaded with the same partial quotient values as

first counters 14, or different counter values may be loaded depending on the application in which the counters are used.

For example, a graphical application may load a value $(a(n)+1)/2$ into the second counters. In this case, no round-off errors result in the continued fraction. To round to the least values, counters 16 may be initialized to 1. To round to the greatest values, counters 16 may be initialized to the corresponding values in counter 14. In order to synchronize to a clock, a clock generator may load different values, which are greater than the partial quotient values.

Circuitry 18 decrements (24) the partial quotient in an initial counter 22 and compares (27) the resulting value to zero. If the resulting value in counter 22 is not zero, circuitry 18 outputs (32) a $g(i)$ value of a_0 .

If the resulting value in counter 22 is zero, circuitry 18 outputs (33) a $g(i)$ value of a_0+1 . In this case, circuitry 18 loads (34) a value to counter 26 from its corresponding second counter 36. Circuitry 18 also increments the counter 38 that immediately precedes counter 26, namely counter 38, and moves on to the next counter 40. Process 10 is then repeated for the next counter 40.

Process 10 has applicability in any field to control values that are based on successive approximations of a rational quantity, such as periodic output. By way of example, process 10 may be used to generate clock signals.

5 The generation of a universal asynchronous receiver-transmitter (UART) clock signal of 14.7456 megahertz (Mhz) from a 2333.3333 Mhz source (provided, e.g., by a crystal oscillator) produces an output clock signal for every 158.23906 input pulses. In this case, process 10 is used as a
10 digital filter. For a $P/Q = 23333333/147456$, the partial quotients are as follows:

[158,4,5,1,1,2,2,1,1,2,2,1,1,14].

15 Process 10 produces the following a_0, a_0+1 values

158	158	158	159	158	158	158	159	158	158	158	159	158	158	158	159	158	158	158	159
158	158	158	158	159	158	158	158	159	158	158	158	159	158	158	158	159	158	158	158
159	158	158	158	159	158	158	158	158	159	158	158	158	159	158	158	158	159	158	158
158	159	158	158	158	159	158	158	158	158	159	158	158	158	159	158	158	158	159	158

20 As shown in Fig. 3, these values result in an output clock signal 46 for every 158 (a_0) or 159 (a_0+1) input pulses of the 2333.3333 MHz source.

Process 10 may also be used in event synthesis
25 applications, such as an asynchronous transfer mode (ATM) scheduler, to schedule simultaneous voice and data traffic. In this regard, an ATM scheduler typically organizes and

controls traffic between different channels sharing a same communication medium. To transmit voice (e.g., a 64 Kilobits/second {Kbits/s} channel) over an asymmetric digital subscriber line (ADSL) (e.g., 796 Kbits/s line) that also transmits data, an ATM Quality of Service (QoS) scheduler handles a traffic ratio of, e.g., 48 payload bytes and 5 header bytes per ATM cell.

In a representative system, the following constraints may apply:

Voice bandwidth requirements:

$64 * 1024 \text{ bits/sec} / (8 * 48 \text{ bytes per cells}) = 170.66$
cells per second.

ADSL bandwidth available (line rate):

$796 * 1024 \text{ bits/sec} / (8 * (5+48) \text{ bytes per cell}) =$
1922.41 cells per second

Bandwidth remaining for consumption by data is

$1922.41 - 170.66 = 1751.75$ cells per second.

An ATM scheduler regulates traffic to provide the correct bandwidth for voice and to ensure that it will be able to send voice cells at the correct rate, which in this instance is one cell of voice after an average of 10.26 data cells per second

(the ratio of data to voice is 1751.75 : 170.66 which approximately equates to 10.26 : 1). For scheduling purposes, the data to voice ratio is calculated without losing the accuracy caused by floating point calculations.

5 The line traffic rate, ltr, is the total available bandwidth of the ADSL line in bits per second. The ltr in this example is:

$$\text{ltr} = (796 * 1024) \text{ bits per second}$$

10 The voice traffic rate, vtr, is the number of bits consumed by voice traffic per second. The vtr in this example is:

$$\text{vtr} = (64 * 1024) * (53 / 48) \text{ bits per second}$$

15 The data traffic rate, dtr, in bits per second uses the remainder of the bandwidth, as follows:

$$\text{dtr} = \text{ltr} - \text{vtr}$$

20
$$\text{dtr} = (796 * 1024) - [(64 * 1024) * (53 / 48)]$$

The ratio data/voice is dtr/vtr is as follows:

$$\frac{(796 * 1024) - [(64 * 1024) * (53 / 48)]}{(64 * 1024) * (53 / 48)}$$

25
$$(64 * 1024) * (53 / 48)$$

which results in the exact ratio 34816/3392 used to obtain the partial quotients for use by process 10. That is,

$$P/Q = 34816/3392 = (10, 3, 1, 3, 1, 2).$$

5

Using those partial quotients, process 10 produces the following a0, a0+1 values

10

10	1	11	1	10	1	10	1	10	1	11	1	10	1	10	1	10	1	11	1	10	1
10	1	11	1	10	1	10	1	10	1	11	1	10	1	10	1	10	1	11	1	10	1
10	1	10	1	11	1	10	1	10	1	10	1	11	1	10	1	10	1	11	1	10	1

Thus, process 10 results in either 10 or 11 ATM non-voice cells being transmitted between every voice cell transmission, denoted by "1".

15

Some standard scheduling algorithms, such as the Generic Cell Rate Algorithm (GCRA) defined by the ATM forum, which repetitively adds a constant, called the theoretical arrival time (TAT), to the cells, introduce a jitter effect. These jitter effects can be reduced by incrementing the TAT with outputs of process 10, as described below.

20

More specifically, the GCRA has a defect linked to a rounding error. Each time an ATM cell arrives, the TAT is incremented by a value "I". This value "I" is the result of a division between an expected traffic rate, a clock driving an external interface, and an internal clock used for time

25

measurements. Whatever these values are, the TAT cannot be a value that can be represented as an integer or a floating point number. The TAT is a result of division, and ignoring a remainder in the division will lead to the following two problems: (1) either "I" is underestimated by a value δ and, after "N" iterations, an accumulated error $N \cdot \delta$ will be greater than "I", thus making the algorithm miss a cell, which may be a non-conforming cell; or (2) either "I" is overestimated by a value δ and, after "M" iterations, an accumulated error $M \cdot \delta$ will be greater than a limit "L", thus making the algorithm classify a cell as non-conforming.

Process 10 can be used to correct the value "I" and then determine a TAT value for a next cell. Because process 10 is not computationally intensive, it can be used in a network processor. Also, because process 10 can be re-initialized without being interrupted (i.e., re-initialized "on-the-fly"), process 10 can be used to compensate for a detected drive of an external clock used for transmission.

By way of example, assume the following, ATM can carry 16515072 cells every 53 seconds. This is the expected traffic rate (ETR). A device's internal clock is the result of a four-divider applied to a 2133.333 MHz clock source, resulting in a clock signal of about 533 Mhz. The internal clock has a

frequency C_i of $2133333/(4 \cdot 1000)$ Mhz. The internal clock will be incremented every $1000/C_i$ nanoseconds (ns).

The ATM transmission system has the following characteristics. The traffic rate is $16515072/53$ cells per ns, and the ratio P/Q to be used in initialization of process 10 is C_i/ETR , or

$$\begin{aligned} & ((2133333 \cdot 1000000)/(4 \cdot 1000))/16515072/53) \\ & = 1570370125/917504 \\ & = 1711.5678 \end{aligned}$$

Thus, there is one ATM cell every 1711 or 1712 internal clock ticks. By choosing to increment the TAT value by 1711, the TAT value will accumulate an error, resulting in ATM cells being considered late in arriving. By choosing to increment the TAT by a value of 1712, the TAT value will accumulate an error, resulting in cells being considered early in arriving. When the error is greater than a limit "L", ATM cells will be considered non-conforming. If the limit is set to 200 clock ticks, the problem may occur as frequently as $200/0.567$ cells, or every 357 ATM cells.

Process 10 may be used to increment the TAT by returning either the 1711 value or the 1712 value. Either of these values may be added to the TAT, thereby reducing the long-term number of early cells and of late cells. That is, because the

same value is not added every time, there is less of a chance of underestimation or overestimation error. In this regard, the continued fraction of 1570370125/917504 is

5 [1711,1,1,3,5,18,1,2,3,2,1,9,4].

This value is used at initialization of process 10. Process 10 then proceeds as described above.

Process 10 may also be used in computer graphics. More specifically, referring to Fig. 4, in computer graphics, a line is drawn by illuminating pixels between two points M,N. To draw a line from the point M to the point N

$$P/Q = (N_x - M_x) / (N_y - M_y) = 10/3$$

15 A continued fraction expansion results in the following continued fraction

$$P/Q = 10/3 = [3,3].$$

20 Applying the partial quotients of the continued fraction to process 10 results in the following output a0,a0+1:

3 4 3 3 4 3 3 4 3 3 4 3 3 4 3 3 4 3 3 4 3 ...

25

As shown in Fig. 5, the values output by process 10 are used to determine a number of (horizontal) pixels 50, 52 to illuminate on a path from point M to point N. Process 10 can be applied to draw lines with a slope greater than 45° by illuminating pixels of a vertical line. Representative computer code that may be used to draw lines in conjunction with process 10 is shown in attached Appendix C.

Process 10 not limited to use with the hardware and software of described herein; it may find applicability in any computing or processing environment

Process 10 can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Process 10 can be implemented as a computer program product or other article of manufacture, e.g., a computer program tangibly embodied in an information carrier, e.g., in a machine-readable storage device or in a propagated signal, for execution by, or to control the operation of, data processing apparatus, e.g., a programmable processor, a computer, or multiple computers. A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for

use in a computing environment. A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

5 Process 10 can be performed by one or more programmable processors executing a computer program to perform functions. Process 10 can also be performed by, and apparatus of the process 10 can be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an
10 ASIC (application-specific integrated circuit).

Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive
15 instructions and data from a read-only memory or a random access memory or both. Elements of a computer include a processor for executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to
20 receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks.

Information carriers suitable for embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in special purpose logic circuitry.

Process 10 can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser, or any combination of such back-end, middleware, or front-end components.

The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network (WAN"), e.g., the Internet.

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The

relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

Representative C++ code to implement process 10 is shown in attached Appendix A; and representative Verilog code to implement process 10 is shown in attached Appendix B.

Other embodiments not described herein are also within the scope of the following claims. For example, the blocks of Fig. 1 may be reordered to achieve the same result. The partial quotients may be generated using any process including, but not limited to, the Euclidian process described in the background, the Greatest Common Denominator (GCD) process, binary reduction, and subtractions.

As noted above, the processes described herein are applicable in a variety of technologies. For example, in addition to the technologies noted above, the processes may be used to distribute P identical objects into Q sets, to arbitrate event collisions, to perform clock correction, and to correct some types of floating-point rounding errors.

What is claimed is: